Branch-and-Price and Constraint Programming for solving a Real-Life Technician Dispatching Problem

Cristián E. Cortés[*]

Civil Engineering Department, Universidad de Chile
Blanco Encalada 2002, 5^{th} floor, Santiago, Chile

Michel Gendreau, Louis Martin Rousseau

CIRRELT and MAGI, Ecole Polytechnique de Montréal
C.P.6079, Succ. Centre-ville, Montreal (Quebec), Canada H3C 3A7

Sebastián Souyris

McCombs School of Business, University of Texas at Austin
2110 Speedway Stop B6500, Austin, TX 78712, USA

Andrés Weintraub

Industrial Engineering Department, Universidad de Chile
Republica 701, Santiago, Chile

Abstract

We consider a real problem faced by a large company providing repair services of office machines~~Xerox~~ in Santiago, Chile. In a typical day about twenty technicians visit seventy customers in a predefined service area in Santiago. We design optimal routes for technicians by considering travel times, soft time windows for technician arrival times at client locations, and fixed repair times. A branch-and-price algorithm was developed, using a constraint branching strategy proposed by Ryan and Foster along with constraint programming in the column generation phase. The column generation takes advantage of the fact that each technician can satisfy no more than five to six service requests per day. Different instances of the problem were solved to optimality in a reasonable computational time, and the results obtained compare favorably with the current practice.

Keywords: Branch-and-price, Constraint programming, Routing, Technician dispatch

problem

[*]Corresponding Author, phone: 562 29784380, fax: 562 26894206

Email addresses: `ccortes@ing.uchile.cl` (Cristián E. Cortés), `michel.gendreau@cirrelt.ca` (Michel Gendreau), `louis-martin.rousseau@polymtl.ca` (Louis Martin Rousseau), `sebastian.souyris@utexas.edu` (Sebastián Souyris), `aweintra@dii.uchile.cl` (Andrés Weintraub)

List of changes

1. Introduction

This research was motivated by a real problem, which is the dispatch of technicians from a large company~~Xerox Chile~~ that provides repair services for their machines to overcome ~~to repair machines~~ failures that occur over a typical working day. The company offers~~Xerox Corporation is an international company that offers color and~~ black-and-white digital printers, digital presses, multifunction devices, and digital copiers. Of the services that the firm provides to its clients, the maintenance of its machines distributed over Santiago is probably one of the most important for building a ~~maintaining the~~ positive image of the company in

2

terms of level and quality of service over time. It is therefore critical to carefully route the available technicians to satisfy the client requirements occurring in certain service areas on a typical working day. In this research, we propose a formulation and solution for this problem, satisfying the typical routing constraints as well as soft time-window constraints. The latter constraints arise from the fact that service requests come from clients with different priorities. The company defines different promised service times (denoted hereinafter target service times) according to the importance of each customer. These promises of technician arrival can be violated, so these conditions can be added as penalties in the objective function. From this, we see that the proposed scheme is based on the classical formulation of the vehicle routing problem with soft time windows (VRPSTW), which is formulated and solved using constraint programming (CP)-based column generation (Section 3).

The vehicle routing problem (VRP) involves constructing routes for a set of vehicles to serve a set of customers, given a number of requirements. There are many classes of the VRP, depending on the requirements and constraints.

One of the best-known versions is the vehicle routing problem with time windows (VRPTW), which has the characteristic that each of the clients must be served within a predefined time interval. VRPTW has been one of the most intensively studied NP-hard problems in the last decades. Exact methods are still restricted to solve a limited number of real instances, and their performance strongly depend on the time-window characteristics. Heuristic and meta-heuristic approaches have been developed for many cases and potential applications of VRPTW. A complete review of the models and methods to solve the VRPTW can be found in Cordeau et al. (2002). Later on, Bräysy and Gendreau (2005a,b) performed deep reviews of VRPTW solution methods, focused on the description of heuristic route constructions methods, local search algorithms and meta-heuristic~~metaheurstic~~ methods for solving capacitated problems of different size and time-window~~s~~ configurations. Nowadays, state of the art results are currently offered by the hybrid genetic algorithm due to Nagata et al. (2010). Their method combines powerful route minimization procedures, proposing an effective edge assembly crossover along with very efficient local search algorithm. In the same line of research, Ibaraki et al. (2008) and Pisinger and Ropke (2007) developed two simple, efficient and flexible methods able to address various VRP variants, namely the

Iterated Local Search (ILS) method and the Adaptive Large Neighbourhood Search (ALNS), respectively. Recently, Vidal et al. (2013) proposed an efficient Hybrid Genetic Search with Advanced Diversity Control for a large class of time-constrained vehicle routing problems, adding to the method new features to properly handle the temporal dimension.

The most successful exact methods are based on column generation for linear programs with decomposable structures; this was originally introduced in Dantzig and Wolfe (1960). The first application of this approach to the VRP was developed by Desrochers et al. (1992) (see also Desrosiers et al., 1995). The basic idea is to decompose the problem into sets of customers visited by the same vehicle (routes) and to select the optimal set of routes among all possible routes. The decomposition is based on two structures: a master problem and a subproblem. The former optimizes the global route objective function and includes the constraint that each customer must be covered exactly once, resulting in a set partitioning formulation. The latter creates new routes that improve the solution, using dynamic programming to solve a shortest path problem with time windows. More recently column generation methods have been extended to include cuts, heuristic pricing and relaxed elementarity constraints (Desaulniers et al., 2008) to produce better lower bounds for the master problem and to speed up the generation of columns. Exact solution approaches using column generation without requiring Branch and Price have been proposed by Baldacci et al. (2012).

With regard to vehicle routing schemes for technician dispatch problems, it is worth mentioning a real application of technicians dispatch for emergency calls of an important electricity company in Chile (Weintraub et al., 1999). Blakeley et al. (2003) also develop and implement a technician dispatch system for a big elevator company using heuristics and GIS data. Recent publications related to service technician routing and scheduling are not many. We can mention the work by Cordeau et al. (2010), who address a service technician scheduling problem arising in large telecommunications companies, focusing on team configurations and assignment of tasks respecting specific skills, priorities and precedence constraints for the tasks assigned, with the objective of minimizing a weighted function related to the makespan. Xu and Chiu (2001) also perform task scheduling for a telecommunication company with the objective of maximizing the realized orders. Tang et al. (2007) implement a

tabu search heuristic for a real world maintenance dispatch system which is formulated as a Multiple Tour Maximum Collection Problem with Time-Dependent rewards. Liberatore et al. (2011) present an exact solution procedure for the VRPSTW. The proposed algorithm is a column generation scheme, with ad-hoc heuristics based on dynamic programming to generate new columns. We use CP for that purpose instead. Liberatore et al. (2011) model explicitly the time window violation for early and late arrivals. By contrast, in our work we just penalize for late arrivals due to the nature of the real application we are trying to solve; nevertheless, in our model and solution algorithm it is straightforward to incorporate penalization for early arrivals as explained later in section 3.2. Kovacs et al. (2012) formulate a service technician routing and scheduling problem motivated by a real problem faced by infrastructure service and maintenance providers. The objective is to minimize the sum of total routing and outsourcing costs, considering hard time windows for reaching customers sites. They solve the problem through an Adaptive Large Neighborhood Search algorithm, tested on both artificial and real-world instances. Pillac et al. (2013) adapt the large neighborhood search algorithm (Shaw, 1998) for the VRPTW with resource constraints. Finally, Souyris et al. (2013) formulate a robust optimization version of a column generation scheme for a technician dispatch problem, putting the focus on properly handling the uncertainty associated with service time in such kind of real applications.

The CP-based column generation (CG) framework was introduced by Junker et al. (1999b) and Yunes et al. (2000) for two different crew rostering problems. Their work was motivated by the difficulties that arose in standard column generation approaches when modeling complex rules from legislation and union agreements. To overcome these difficulties, they proposed solving the pricing subproblem using CP models of resource-constrained shortest paths on acyclic graphs. The most attractive feature of CP, compared with dynamic programming, was the expressiveness of its modeling languages. Rousseau et al. (2004) and then Chabrier (2006) extended this framework to cyclic graphs, thus allowing VRPTW problems to be solved. In the last decade, the CP-CG framework has been used in several applications that are discussed in a recent survey by Gualandi and Malucelli (2009).

In this paper we use a CP-based pricing model similar to that proposed by Yunes et al. (2000, 2005) in the context of bus driver scheduling, which allows for a flexible and

simple modeling of the technician-route requirements. This straightforward model, which takes advantage of the fact that each technician visits only a small fraction of the overall daily clients, has significantly fewer variables than the graph-inspired models normally used for routing problems (Rousseau et al., 2004). Although it may be somewhat less efficient, this model does not require the implementation of complex shortest-path constraints such as those proposed for airline crew rostering problems (Junker et al., 1999b). To obtain optimal solutions for different service request distributions, we implement a branch-and-price approach based on the efficient branching strategy of Ryan and Foster (1981), again taking advantage of CP flexibility when we include these branching constraints.

~~The main contributions of this paper are the following: (i) we provide what we believe is the first exact solution procedure for the VRPSTW; (ii) it is also the first to show that the approach proposed by~~ Yunes et al. (2000, 2005) ~~for crew scheduling problems can be applied to problems whose underlying structure is a cyclic graph, and to integrate it into a branch-and-price framework; (iii) as stated above, the approach that we propose has some advantages over the alternatives; it can thus prove attractive for similar practical applications.~~

The main contributions of this paper are three-fold: (i) we provide an efficient exact solution procedure for the VRPSTW using CP; (ii) it is the first to show that the approach proposed by Yunes et al. (2000, 2005) for crew scheduling problems can be applied to problems whose underlying structure is a cyclic graph, and to integrate it into a branch-and-price framework; (ii) as stated above, the approach that we propose has some advantages over the alternatives; it can thus prove attractive for similar practical applications.

The remainder of this paper is organized as follows. In the next section we state the problem formally and provide a mathematical formulation. Section 3 presents the column generation approach,~~for both the master and the pricing problem,~~ while the branch-and-price methodology is described in Section 4. Experiments on a real-case scenario are reported in Section 5.

2. Problem Statement and Mathematical Model

The strategic objective of the firm is based on client satisfaction. Within this context, the maintenance of their machines is one of the most important activities of the company in Chile. As explained in Section 1, service requests have different priorities, and there are different target response times for service requests at different priority levels. Strictly, the target response time is defined as the maximum allowable time for a technician to reach the service location, measured from the time of the service request. If the technician reaches the location after the target response time, a penalty will be incurred by the system; this is considered in the formulation.

To consider this effect, we use a compound objective function for assigning technicians and jobs. The function minimizes two components: the sum of the differences between the target response times of requests and the effective service times provided by the firm, plus the travel times. This objective function seems to meet the needs of the company since it takes into account both service quality and the effective use of technicians.

The approach considers expected travel and service times, which are estimated from historical company data. In further developments, we will include uncertainty in the service times to make the model more robust and realistic. Travel times are less important than service times in the results, and therefore the assumption of deterministic travel behavior is reasonable in this case.

The set of service requests assigned during a given day come from the previous days, usually the day before, since the company attempts to enforce a 24-hour service policy. It is assumed that the dispatcher in charge of service selects which requests should be handled during the coming day. Furthermore, the dispatcher will choose a set of high-priority requests to be served first. A technician will begin his working day at one of these high-priority customer locations, under the reasonable assumption that the number of service requests is larger than the number of available technicians. In some cases, a technician can start the working day at the depot because he was not assigned to a specific client. In addition, all the tasks assigned to each technician must be completed during the day.

The modeling scheme was adapted from the classical formulation of the VRPSTW

(Cordeau et al., 2002) with the following adjustments:

- For this problem, only the upper bound of the soft time window (related to the target response time of each service request) is considered, since the objective is to serve each requirement as soon as possible.

- Service times are quantitatively longer and are known less accurately than travel times; therefore, the former will play a more important role than the latter in the proposed decision rules.

- Although the VRPSTW considers soft time windows by penalizing the objective function, we add the hard constraint that all tasks scheduled for a specific day must be completed during that day. To ensure feasibility, we modified the classical formulation slightly to be able to decide which service requests can be handled during the day by the available technicians; any remaining requests are then postponed to the next day and given a very high priority.

Let $\mathcal{K} = \{1, \ldots, K\}$ be the set of technicians available to work during the day. As mentioned earlier, the technicians must start the day at the location of a service request with a very high priority or at the depot. Let $I_1 = \{1, \ldots, K\}$ be the set of these locations; and $I_2 = \{K + 1, \ldots, C\}$ be the set of locations of the service requests that remain to be scheduled, where C is the cardinality of $I_1 \cup I_2$. We set $C + 1$ as a dummy depot to which all technicians must be sent after their schedule is finished. To simplify the notation, we define the arc set $A = \{(i, j) \mid i \in I_1 \cup I_2, j \in I_2 \cup \{C + 1\}, \ i \neq j\}$ to represent all feasible trips between locations. Additionally, let b_i be the time-window upper bound of service request i, and s_i the expected repair time for service request $i, i \in I_1 \cup I_2$. Let t_{ij} be the travel time from service request i to service request j, $(i, j) \in A$. We assume that travel times satisfy the triangular inequality. The end of the working day is set at instant F, which represents the latest time a technician can start the last job of the day. $\beta \in [0, 1]$ is a multiobjective parameter.

We next formulate a mixed integer model, including both $\{0\text{-}1\}$ (binary) and continuous variables. Specifically, we define flow variable x_{ij}^k, which is equal to 1 if technician k attends service request i and service request j sequentially, and 0 otherwise, $(i, j) \in A, k \in \mathcal{K}$.

Two continuous variables are also included: w_{ik} is the time that technician k arrives at the location of service request i, $i \in I_1 \cup I_2 \cup \{C+1\}, k \in \mathcal{K}$; and δ_{ik} represents the amount of time by which technician k violates the soft time window of service request i, $i \in I_2, k \in \mathcal{K}$. Finally, we define the binary variable v_i, which allows the model to handle the demand not served during the day with the available fleet by scheduling those service requests for the beginning of the next day, $i \in I_2$. In the model, such requests are assigned to virtual paths at a high penalty P. Thus, v_i is equal to 1 if i is sent through a virtual path, 0 otherwise. To simplify the notation we define the sets $O(i)$ and $D(i)$ associated with node i: $O(i) = \{j \mid (i,j) \in A\}, i \in I_1 \cup I_2$, and $D(i) = \{j \mid (j,i) \in A\}, i \in I_2 \cup \{C+1\}$. The formulation is as follows:

$$\min_{x,w,v,\delta} \quad \beta \sum_{k \in \mathcal{K}} \sum_{i \in I_2} \delta_{ik} + (1-\beta) \sum_{k \in \mathcal{K}} \sum_{(i,j) \in A} t_{ij} x_{ij}^k + \sum_{i \in I_2} P v_i \tag{1}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}} \sum_{j \in O(i)} x_{ij}^k = 1 \qquad\qquad i \in I_1 \quad (2)$$

$$\sum_{k \in \mathcal{K}} \sum_{j \in O(i)} x_{ij}^k = 1 - v_i \qquad\qquad i \in I_2 \quad (3)$$

$$\sum_{j \in O(i)} x_{ij}^k - \sum_{j \in D(i)} x_{ji}^k = \begin{cases} 1 & i = k \\ -1 & i = C+1 \\ 0 & \text{otherwise.} \end{cases} \qquad i \in I_1 \cup I_2 \cup \{C+1\}, \ k \in \mathcal{K} \quad (4)$$

$$w_{ik} + s_i + t_{ij} - w_{jk} \leq (1 - x_{ij}^k) F \qquad\qquad (i,j) \in A, \ k \in \mathcal{K} \quad (5)$$

$$w_{ik} \leq F \sum_{j \in D(i)} x_{ji}^k \qquad\qquad i \in I_2, \ k \in \mathcal{K} \quad (6)$$

$$w_{ik} - \delta_{ik} \leq b_i \qquad\qquad i \in I_2, \ k \in \mathcal{K} \quad (7)$$

$$x_{ij}^k \in \{0,1\} \qquad\qquad (i,j) \in A, \ k \in \mathcal{K} \quad (8)$$

$$v_i \in \{0,1\} \qquad\qquad i \in I_2 \quad (9)$$

$$w_{ik}, \delta_{ik} \geq 0 \qquad\qquad i \in I_2, \ k \in \mathcal{K}. \quad (10)$$

The objective function (1) accounts for the total cost, computed as a convex combination of the sum of soft time-window violations and the total travel time. An additional term is added to penalize the unsatisfied demand. Constraints (2) and (3) restrict the assignment of each

service request scheduled during the day to exactly one technician; if the request is postponed to the next day, the condition $v_i = 1$ together with constraint (3) ensure that no technician attends it on the current day. Next, constraints (4) ensure flow conservation, considering the dummy node where each path must end. Constraints (5) guarantee schedule feasibility with respect to precedence-time consistency (i.e., if technician k travels from customer i to customer j, then the arrival time at node j will be greater than or equal to the arrival time at node i plus the service time there). If $x_{ij}^k = 1$, then k goes from i to j, and we can start serving j no earlier than the start of service at i, plus the time spent in i, plus the travel time from i to j. If $x_{ij}^k = 0$, technician k does not go from i to j and constraint (5) becomes inactive. Note that, for a given k, constraints (6) force $w_{ik} = 0$ if service request i is not met by technician k and require k to reach i before time F if service is going to be provided (i.e., $v_i = 0$). Constraints (7) define δ_{ik}, the violation of the soft time window. Finally, (8), (9), and (10) impose binary conditions on the flow variables and state the nonnegativity restriction on the time-arrival variables.

The VRP problem with soft time windows is hard to solve as mentioned in Taillard et al. (1997). We propose to solve this problem using column generation, as described in the next section.

3. Column Generation Approach

By working with actual company data, we realized that most real instances were not solvable using the arc-based formulation presented in the previous section, because the number of variables was too large to handle. Therefore, we looked for an alternative formulation and solution approach.

Column generation approaches have provided promising results for various types of VRPs (see for example Desrochers et al., 1992). Moreover, column generation approaches are very flexible in the sense that the problem under consideration can be split into two parts: a main model, the master problem, which chooses the routes with the minimum total cost from a pool of feasible routes; and a secondary model, the subproblem, which generates feasible routes that could potentially reduce the total cost.

In the literature, the subproblem is usually solved via dynamic programming (DP). We decided to use CP instead, for several reasons. First, it seems that CP works well when the length of each route is relatively small, which is the case for this application. Second, the CP model we implement is simple and easy to code. Third, under a CP approach, any additional constraint can be incorporated directly into the code without a special modeling technique. This is crucial when we use the branch-and-price algorithm (see Section 4) improved by the branching strategy proposed by Ryan and Foster (1981). Imposing some constraints on routes turned out to be straightforward with CP. For example, the proposed branching strategy forces two specific service requests to appear in a specific route, which can be included easily under CP, whereas in a DP implementation incorporating such constraints could become cumbersome. In Sections 3.2 and 4 we describe these features in detail.

3.1. Master Problem

The master problem can be formulated as a set partitioning model assuming that it is possible to choose routes for each technician from an existing set of routes R. We use the result presented in Barnhart et al. (1998), which shows that the set covering relaxation is numerically far more stable and thus easier to solve than the set partitioning version. Also the authors argue that it is easy to construct a feasible integer solution from a solution of the set covering relaxation.

Each route $r \in R$ is characterized by a technician who initiates the route r at a specific service location $i_1 \in I_1$, and then follows a sequence of service requests $\{i_2, \ldots, i_e\} \subseteq I_2$, where i_2 is the second one and i_e is the service request at the end of route r. For each service request in position $l \in \{2, \ldots, e\}$ of the route, the technician's arrival time is $w_{i_l} = w_{i_{l-1}} + s_{i_{l-1}} + t_{i_{l-1}i_l}$, and thus the time-window violation becomes $d_{i_l} = \max\{0, w_{i_l} - b_{i_l}\}$. Additionally, $w_{i_1} = d_{i_1} = 0$. Hence, the total cost of the route is $c_r = \beta \sum_{l=2}^{e} d_{i_l} + (1 - \beta) \sum_{l=1}^{e-1} t_{i_l,i_{l+1}}$. In this formulation, the binary variable θ_r indicates whether or not route $r \in R$ is chosen. a_{ir} is a binary parameter that indicates whether or not route r contains service request i, and the binary variable v_i is equal to 1 if service request i is not included

in any route, incurring a high penalty P. Then the master problem is the following:

$$(\text{MP}) \quad \min_{\theta, v} \quad \sum_{r \in R} c_r \theta_r + \sum_{i \in I_1 \cup I_2} P v_i \tag{11}$$

$$\text{s.t.} \quad \sum_{r \in R} a_{ir} \theta_r + v_i \geq 1 \quad i \in I_1 \cup I_2 \tag{12}$$

$$\theta_r \in \{0, 1\} \quad r \in R \tag{13}$$

$$v_i \in \{0, 1\} \quad i \in I_1 \cup I_2. \tag{14}$$

Note that in the MP formulation, the pool R of routes can be empty, and there still exists a feasible solution: $v_i = 1, i \in I_1 \cup I_2$, with cost $C \times P \dot{C} \dot{P}$, where C is the cardinality of $I_1 \cup I_2$ as defined before.

To generate new columns, we replace the integrality constraints (13) and (14) by $\theta_r \in [0, 1], r \in R$, and $v_i \in [0, 1], i \in I_1 \cup I_2$, respectively. Thus, for a given pool of routes R, the optimal solution of MP provides the dual values of constraint (12) to the subproblem that generates new routes. This procedure is described below.

3.2. Subproblem

Given a pool R of columns and the associated optimal solution of the LP relaxation of problem (11) to (14), it is well known from linear programming theory that a new column r, not in R, has the potential to improve the objective function only if it has a negative reduced cost. The reduced cost of a column is defined as the cost of the route, c_r, minus the sum of the dual variables of the service requests that belong to that route, which come from constraint (12).

At each iteration, the subproblem identifies a route of minimum reduced cost. We solve the master problem over the current set of columns using the simplex method and obtain the dual variables. Let α_i be the dual variables associated with constraint (12). The subproblem generates the optimal route by minimizing the real cost c_r (the sum of the time-window violations and travel times) minus the sum of the dual variables of the service requests included in the route, subject to the constraints that ensure that the route is feasible. By this process, a route with minimum reduced cost is generated. If the resulting reduced cost

is negative, this new column can be included in the basis to improve the objective function of the master problem.

Constraint programming models have been used to solve pricing subproblems within hybrid column generation for almost a decade (Gualandi and Malucelli, 2009). There are essentially three families of models in the literature that specifically address routing and scheduling problems. One of the first models, proposed by Yunes et al. (2000, 2005), is based on an array of finite-domain variables $X_p \in T, p \in P$, that identify which task in T is to be performed by a bus driver in position $p \in P$. Junker et al. (1999a,b) and Fahle et al. (2002) proposed using a single set variable $S \subseteq T$ to identify the subset of tasks in T to be covered by a crew pairing. In order for their model to be valid and efficient, they introduced a new shortest path constraint, as well as a negative-reduced-cost constraint on S. Finally, in the context of vehicle routing with time windows, Rousseau et al. (2001, 2004) based their model on successor variables $N_t \in T, t \in T$, that identify the task to be performed immediately after t. This approach also required the use of specially designed global constraints and a search strategy based on dynamic programming.

In this paper, we adapt the simple model proposed by Yunes et al. (2000) for the following reasons. First of all, it is straightforward, is flexible, and works without the addition of dedicated global constraints. In constraint programming, a global constraint is a relationship among decision variables for which an efficient algorithm is available that can find the set of all infeasible values for each of the included decision variables. For example, the global constraint alldifferent(x_1, x_2, x_3) ensures that the three variables x_1, x_2, and x_3 are all different. Modern constraint programming packages offer a set of global constraints, and associated algorithms, that can model a generous number of relationships. The algorithms to find feasible sets are based on the principles of constraint propagation and domain reduction. When it is not possible to model a particular relationship with the existing global constraints, it is necessary to design a dedicated global constraint and an associated algorithm. A detailed explanation of these concepts is found in the exhaustive survey by Régin (2011).

Second, in contrast with the model of Rousseau et al. (2004), which defines one variable for every node or task in the problem, Yunes et al. (2000) introduce only a number

of variables equivalent to the number of tasks that can be performed by one technician. Since in our context this number is small (from three to six), this model seems particularly appropriate. Moreover in both these models the size of domains of the main decision variables are equal as they correspond to the number of nodes in the graph.

In the subsequent model (15) to (25), the new route to be generated is represented by the array of variables $sc[l], l = 1, \ldots, L$, where L is the maximum number of service requests that a technician can satisfy in one day. The l^{th} element of sc is the service request scheduled in position l of the route. The route must start at an initial location of a technician (which could be either a high-priority service request or the depot, in case there are more available technicians than high-priority requests), that is a location from I_1. The route must contain at least one service request that remains to be scheduled, that is a service request from set I_2. Any position of the route that is not used by a service request must be utilized by fictitious nodes that are added for CP modeling purposes. Hence, the maximum number of fictitious nodes that a route can contain is $L - 2$. Let us define the set of fictitious nodes to be $I_3 = \{C + 1, \ldots, C + L - 2\}$, where, as before, C is the cardinality of $I_1 \cup I_2$. Therefore, the domain of the $sc[l]$ variables, $l = 1, \ldots, L$, is set to $I_1 \cup I_2 \cup I_3$. The distinction between the fictitious nodes in I_3 is necessary in order to impose the global constraint alldifferent(sc), which will be discussed after the presentation of the model. The travel time between a service request in I_2 and any fictitious node in I_3 is set to 0, as is the travel time between any pair of nodes in I_3. Figure 1 shows an example of a route that starts servicing request 3, then proceeds to requests 5, 10, and 14, and finishes with the fictitious nodes $C + 1$ and $C + 2$.

Variables $w[l]$ and $d[l]$ define the start time of service and the time-window violation respectively for the service request in position $l = 1, \ldots, L$. Variable $a[i]$, $i \in I_1 \cup I_2 \cup I_3$ takes the value 1 if request i is served by the route, and 0 otherwise, $l = 1, \ldots, L$. These definitions are useful for the control of the branch-and-price method as described in Section
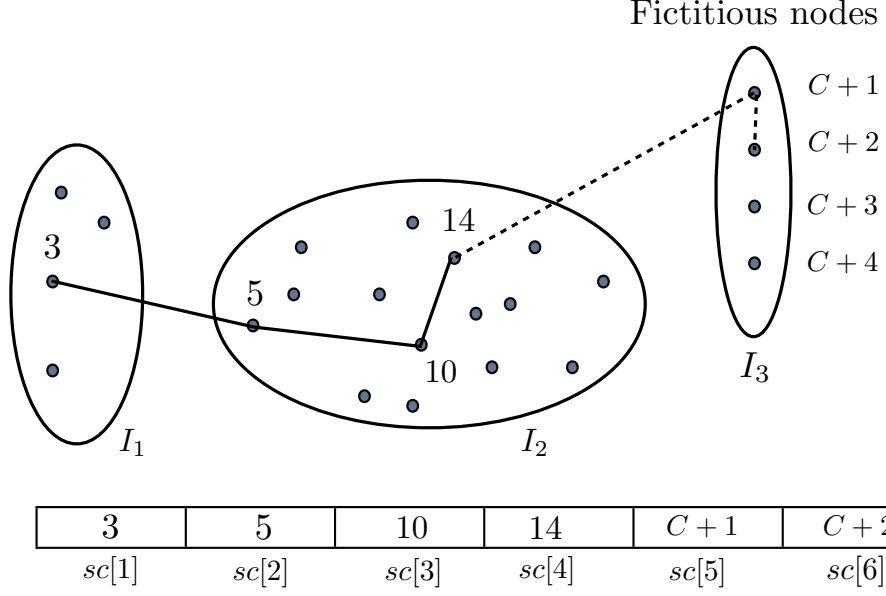
Figure 1: Shortest path model implemented with CP

4. The constraint programming model for the subproblem is thus the following:

$$\text{(SP)} \quad \min \beta \sum_{l=1}^{L} d[l] + (1-\beta) \sum_{l=1}^{L} t_{sc[l-1],sc[l]} - \sum_{l=1}^{L} \alpha_{sc[l]} \tag{15}$$

$$\text{s.t. } w[1] = 0 \tag{16}$$

$$w[l] = w[l-1] + s_{sc[l-1]} + t_{sc[l-1],sc[l]} \qquad l = 2,\ldots,L \tag{17}$$

$$d[l] = \max(0, w[l] - b_{sc[l]}) \qquad l = 1,\ldots,L \tag{18}$$

$$\text{alldifferent}(sc) \tag{19}$$

$$sc[l] \in I_1 \qquad l = 1 \tag{20}$$

$$sc[l] \in I_2 \qquad l = 2 \tag{21}$$

$$sc[l] \in I_2 \cup I_3 \qquad l = 3,\ldots,L \tag{22}$$

$$sc[l] = i, i \in I_3 \Rightarrow sc[l+1] = i+1 \qquad l = 3,\ldots,L \tag{23}$$

$$a[sc[l]] = 1 \qquad l = 1,\ldots,L \tag{24}$$

$$\sum_{i \in I_1 \cup I_2 \cup I_3} a[i] = L \tag{25}$$

The objective function (15) minimizes the convex combination of time-window violation and

15

travel cost, which is the real cost of the route minus the sum of the dual values, α, associated with the service requests in the route. For $i \in I_3$, which is not in constraint (12) of MP, we set α_i equal to 0. Constraint (16) ensures that the start time of the path is associated with the start time of the day, and (17) sets the starting time of next service requests that are visited by the route. CP constraint (18) sets the late time-window violations. Note that the model can also include early time window violations; for adding that feature, we just have to change the right hand side of constraint (18) by $\max(f_{sc[l]} - w[l], w[l] - b_{sc[l]})$, where f_i corresponds to the time window lower bound associated with service request i. Thus, the CP algorithm presented next can be easily adapted to incorporate this modification. Constraint (19), alldifferent(sc), is a global constraint that ensures that all of the variables sc are different, so any request is served by the path no more than once (for an explanation of this global constraint see, for example, Régin (2011)). Constraint (20) ensures that the first request served by the route is assigned to a location where a technician is initially positioned, (21) ensures that the second position of the route is a service request without a technician initially allocated, and (22) restricts the other positions along the path to be either service requests or fictitious nodes. (23) ensures that, if position l is used by a fictitious node, the next service request in the route must be the next fictitious node (proceeding to any other service request is not allowed in order to reduce the search space as much as possible). Variable a is useful to impose new constraints for the service requests that can belong to a specific route, which are needed to implement the search procedure in the branch-and-price methodology, as explained below and detailed in Section 4. Finally, constraints (24) and (25) impose logical restrictions between the different variables in sc and a.

CG frameworks depend heavily on marginal costs to guide the search at the subproblem level. In some cases it is possible that, during the first iterations, the marginal cost associated with each customer is not accurately estimated by the dual values. For instance, in some routes some service requests pick up most of the total dual values. This undesirable behavior is illustrated by the example in Fig. 2. A path that visits each of overweighted service requests (2 and 3's dual values are exactly twice 1 and 4's dual values) will be considered a good route (with a reduced cost of -5), but it is not (it is unlikely to be selected in an optimal IP solution). With a more realistic distribution of the dual values, all the nodes

would have been given a value of 15 and no more reduced-cost paths would have been found, thus removing the need for a last iteration.
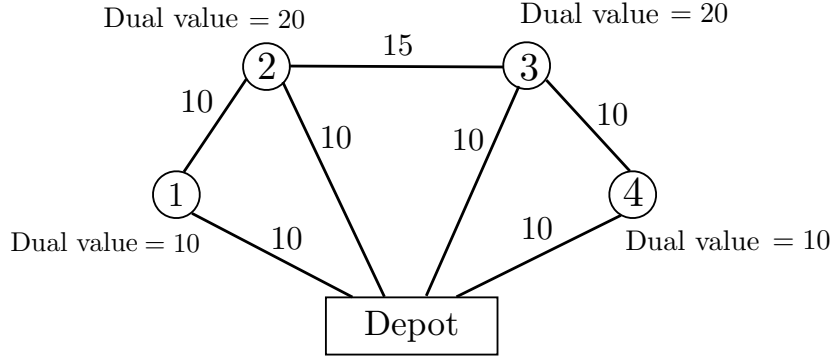

Figure 2: Path $Depot - 2 - 3 - Depot$ has a negative reduced cost of -5.

In this example, described in Rousseau et al. (2007), such undesirable behavior may occur because the master problem is degenerate, and thus its dual has an infinite number of optimal solutions. The bases of the primal solutions all exhibit a common set of strictly positive variables, but different sets of null variables. A different dual solution is associated with each of these equivalent primal optimal solutions.

The standard function that returns dual values in the LP codes returns an extreme point of the dual polyhedron. Extreme solutions are characterized by large values for some marginal costs while others are at zero. Therefore, the subproblem tends to build routes that have very low reduced costs but potentially large travel times and time-window violations. To avoid this, we instead optimize over the real objective function (c_r). Negativity of the reduced cost can then be enforced through a constraint, since the necessary condition imposed by the CG framework to ensure convergence is that we add, at each iteration, a set of negative-reduced-cost routes or prove that none exists. The objective function (15) thus becomes (26) and constraint (27) is added to the subproblem model:

$$\min \beta \sum_{l=1}^{L} d[l] + (1 - \beta) \sum_{l=1}^{L} t_{sc[l-1],sc[l]} \tag{26}$$

$$\beta \sum_{l=1}^{L} d[l] + (1 - \beta) \sum_{l=1}^{L} t_{sc[l-1],sc[l]} - \sum_{l=1}^{L} \alpha_{sc[l]} < 0 \tag{27}$$

After experimenting with our CP implementation, we decided to split the subproblem into a two-step procedure (SP1 and SP2); we found a considerable reduction in the running time when many interesting columns are generated at each iteration of the pricing problem. SP1 attempts to select service requests such that the reduced cost is negative (i.e., constraint (27) is satisfied), but the order of the service requests in the route is not necessarily the best. Thus, SP1 is not an optimization problem but a CP feasibility problem with feasible region (17) to (27). SP2 then optimizes the sequence to find the lowest real cost given by (26) subject to (17) to (25). SP1 passes the selected service request to SP2 by fixing the variables a in constraint (24). Figure 3 shows an example of the process. First, SP1 chooses service requests 0, 2, 3, 5, 8, and 10, and then SP2 rebuilds the route with only these service requests. Notice that the reduced cost of the resequenced route can only be lower than the original route or equal in case the sequence does not change, because the real cost of the re-sequenced route decreases and the sum of the dual variables remains the same.

$sc[1] = 0;\ sc[2] = 2;$
$sc[3] = 3;\ sc[4] = 5;$ $\Rightarrow$ $a[0] = a[2] = a[3] = a[5] = a[8] = a[10] = 1;$ Reconstruct the path
$sc[5] = 8;\ sc[6] = 10;$ $\quad a[4] = a[7] = a[9] = a[11] = a[12] = a[13] = 0;$ with the nodes.

$sc[1] = 0;\ sc[2] = 3;$
$sc[3] = 8;\ sc[4] = 5;$
$sc[5] = 2;\ sc[6] = 10;$

Figure 3: Construction of a new route procedure.

Empirically, we found a considerable improvement in terms of the number of columns generated using this procedure, which reduces the computational time by about 30%.

To help SP1 find routes of good quality faster, we implement a simple but efficient search consistent with our problem. It guides the local search logically, by exploring a neighborhood matrix for each service request.

For a given set of dual values α, we define the neighborhood matrix $N(\alpha)$ to be a square matrix of C rows and columns (recall that $C = |I_1 \cup I_2|$). The element in row i and

column p, n_{ip}, is the p^{th} closest service request to request i. We considered different distance measures, and the one that performs the fastest in computational terms is the following. For a new route that has been created and includes the service request i, we would like to have the sequence (i, j) if two conditions hold: first, j's time-window upper bound, b_j, is close to $b_i + s_i + t_{ij}$ (recall that s_i is the expected service time of i, and t_{ij} is the expected travel time between i and j); and second, j has a large dual value. If j has a small dual value, then it is already efficiently covered by a route belonging to pool R. Thus, for each pair (i, j) we define the distance measure $d(i, j) = \gamma((b_j - (b_i + s_i + t_{ij}))^+ + \rho(b_i + s_i + t_{ij} - b_j)^+) + (1 - \gamma)\frac{1}{\alpha_j}$, where $(x)^+ = \max\{0, x\}, \rho > 1$ and $\gamma \in [0, 1]$. The larger penalization when $b_i + s_i + t_{ij} > b_j$ is to avoid a time-window violation at j. The implemented search for SP1 that uses $N(\alpha)$ is described in Algorithm 1.

Algorithm 1: SP1 constructs a new route with negative reduced cost, but not necessarily in the optimal order

---

Input : Dual values α from the master problem. Neighborhood matrix $N(\alpha)$
Output: A new route with negative reduced cost but not necessarily in optimal order

---

1 $sc[1] \leftarrow \arg\max_{i \in I_1}\{\alpha_i\}$;
2 $a[sc[1]] \leftarrow 1$;
3 **for** $l \leftarrow 2$ **to** L **do**
4 $i \leftarrow sc[l - 1]$;
5 $p \leftarrow 1$;
6   **while** $sc[l]$ *has no service request assigned to it* **do**
7   **if** *it is feasible by (17) to (27) to assign service request n_{ip} to $sc[l]$* **then**
8 $sc[l] \leftarrow n_{ip}$;
9 $a[sc[l]] \leftarrow 1$;
10   **else**
11 $p \leftarrow p + 1$;

We also notice an additional improvement in the performance of the subproblem when adding some redundant constraints that do not modify the solution set but improve the constraint propagation. These constraints improve the tree pruning, filtering processes, and domain reduction. We impose

$$sc[l] > first(I_3) \Rightarrow sc[l - 1] = sc[l] - 1 \qquad l = 3, \ldots, L \qquad (28)$$

$$sc[l] \leq first(I_3) \Rightarrow sc[l - 1] < first(I_3) \qquad l = 3, \ldots, L \qquad (29)$$

where $first(I_3)$ represents the first fictitious service request; i.e., $C + 1$. Thus, constraints (28) and (29) eliminate identical solutions from the domain, taking into account the fact that all the fictitious nodes are the same, although their identification codes are different. Without these constraints, the solver does not realize early in the search tree that a service request cannot follow a fictitious node. In the CP literature, redundant constraints have often been used to create a set of different links between variables and to generate new potential for propagation.

In the implementation, instead of adding only one route per iteration of the column generation, we generate a number of routes with negative reduced costs. To do this, we take advantage of the CP search tree of SP1. Each final node of the search tree is a route solution that has a negative reduced cost; therefore, each has the potential to improve the solution of MP if it is added to the pool R. We add the first Q columns found. Between iterations, we tune Q depending on the time that SP1 is taking to find solutions. For the first iterations we set Q to a large number, and for the last iterations we set Q to 1.

We implemented a simple strategy to control the size of the pool of columns. At every iteration of the CG process, columns with high positive reduced costs are eliminated from the pool.

Finally, we implemented a branch-and-price method that guarantees finding the optimal solution. The method allows us to explore additional routes with negative reduced costs in the branch-and-bound tree, by inspecting each node of the tree rather than only the root node. The method performs significantly better when an ad-hoc branching strategy is implemented, as explained below.

4. Branch-and-Price Implementation

Many successful implementations of branch-and-price schemes can be found in the literature. Barnhart et al. (2000) presented a CG model that is solved with a branch-and-price-and-cut algorithm for an origin-destination integer multicommodity flow problem. Savelsbergh and Sol (1991) solved a dynamic vehicle routing problem with heuristic optimization techniques based on dynamic programming, together with a sophisticated column man-

agement scheme. Theoretical aspects of several versions of the branch-and-price algorithm are discussed by Barnhart et al. (1998), Vanderbeck (2000, 2005, 2006), and Vanderbeck and Savelsbergh (2006). Specifically, the authors concentrate on topics such as the formulation of the decomposition, proper ways to perform the branching, and column-management efficiency.

In this application, we use the branching strategy proposed by Ryan and Foster (1981) for generic set partitioning problems. Barnhart et al. (1998) proved the following proposition:

Proposition: If A is a 0-1 matrix, and a basic solution of $A\theta = 1$ is fractional i.e., at least one of the components of θ is fractional, then there exist two rows i and j in the master problem such that:

$$0 < \sum_{r:a_{ir}=1,a_{jr}=1} \theta_r < 1 \qquad (30)$$

The pair i, j establishes the following pair of branching constraints: $\sum_{r:a_{ir}=1,a_{jr}=1} \theta_r = 1$ and $\sum_{r:a_{ir}=1,a_{jr}=1} \theta_r = 0$, i.e., rows i and j must be covered by the same column on the first (left) branch and by different columns on the second (right) branch. The CP model can incorporate constraints of this type in an efficient way.

Thus, for the left branch, constraint (31) must be added to SP1 to restrict the columns to those containing either both service requests i and j, or neither of them:

$$a[i] + a[j] \neq 1 \qquad (31)$$

Similarly, for the right branch, constraint (32) can be added to restrict the columns to those containing at most one of the service requests, either i or j:

$$a[i] + a[j] \leq 1 \qquad (32)$$

Many possible pairs of service requests (i, j) could be chosen. For efficiency, we use a rule similar to that proposed by Vance et al. (1997) to choose a pair of service requests with a large probability of being covered by the same route in a good feasible solution of the IP. We compute for each pair (i, j) the scalar $f(i, j) = \sum_{r:a_{ir}=1,a_{jr}=1} \theta_r$. We choose the pair

with the largest $f(i, j)$ for branching, and the depth-first search strategy is applied along the left branch, ensuring that the chosen pair of service requests will be on the same route. As mentioned before, the structure of the proposed CP model allows us to implement the branching strategy in a straightforward way.

Our branch-and-price procedure is terminated when either we find an acceptable gap or the running time exceeds three hours. The gap is computed as the ratio of the integral objective value minus the linear relaxation objective value, and the linear relaxation objective value ((upper bound − lower bound)/ lower bound).

In the next section, we report some empirical results for different service-request sets, using real data for a typical day.

5. Implementation and Empirical Results

The model was coded in Ilog Concert Technology and solved using CPLEX 9.0 (Ilog, 2003a) for the master problem and SOLVER 6.0 (Ilog, 2003b) for the subproblems. We test our algorithms in an Intel Pentium M 1.5 GHz, with 2 GB of RAM. In this section, the branch-and bound-method (B&B), the optimal branch-and-price scheme (B&P), and the linear relaxation at the root node are run for several real instances of different sizes and service-request configurations, all from real data provided by the firmXerox.

The consistency of the CG approach was empirically checked by observing the convergence pattern toward the optimal solution in small problems. In fact, for small instances the optimal solution was obtained directly by solving the IP corresponding to the original arc-based formulation presented in Eqs. (1)–(10).

In our results, each technician was assigned to no more than six service requests per day. For a larger number of requests, it was necessary to reschedule some for the next day, since serving all the requests on the same day was not feasible.

5.1. Description of the experiments with real data

We used a data set for the southern region of Santiago, Chile. Let us denote this region AB, since it can be further split into two subregions, namely A and B. A normal day was considered in terms of the number of service requests and the number of technicians working. Thus, 41 calls were received in zone A with 10 technicians available, while 35 were received in zone B with 9 technicians available. Note that the company presently schedules service for the two subregions separately. In our tests, we consider three basic instances, one for region A, another for region B, and one for the entire region AB. Figure 4 shows a map of the city and the two zones.

Figure 4: Santiago city with zones A and B.

The travel times between service-request locations were estimated by dividing the complete region AB into microzones. We decided to use the administrative divisions of Santiago (the microzones then match what are denoted *comunas* by the Chilean authorities). This allows us to use historical information for average travel times between each pair of comunas and within each comuna. Microzones (comunas) are required because the travel times between service requests are not available. The expected service time is computed as the average service time reported by the technicians for each type of request over the year. Finally, target response times (as defined in the first paragraph of Section 2) were also provided by the company, depending on the priority assigned to each service request.

23

In most cases, the observed travel times followed the triangular inequality. If they did not, we perturbed them with the smallest possible amount necessary to satisfy the triangular inequality. As mentioned earlier, travel times are much smaller than service times, therefore this perturbation does not significantly change the solution.

We considered $\beta = 0.3, 0.6, 0.9$, which should cover the different priority levels a modeler would assign to the two components of the objective function (travel times versus time-window violations) according to the company's goals. In this scheme, if it is infeasible to insert a specific service request into any route, that request will be assigned for service at the beginning of the next day, with a cost-function penalty of 500. Thus, the first columns generated for each service request are conceptually similar to the virtual paths defined in the original model through the variables v_i to handle infeasibility. They are assigned a high penalty to provide an initial feasible solution to the master problem. Under these conditions, with the original IP model, Eqs. (1)–(10), we obtain the same optimal solution as that obtained by B&P. Unfortunately, this result can be verified only for very small instances, since the original arc-based formulation becomes intractable as the number of requests increases.

Under these conditions, we run our model using the B&P algorithm separately for each subregion (A and B), in order to compare our results with the manual dispatch used by the company (company operation) for each subregion. Table 1 compares the objective values of the results. Table 2 compares the results in terms of time-window violation and travel time for the entire system. In Table 2, column Δ reports the total time-window violation observed for all service requests on all prescribed routes, and column T reports the sum of the travel time over all vehicle routes.

Instance A	Objective Function	
β	Manual	B&P
0.3	813	613
0.6	854	578
0.9	896	491

Instance B	Objective Function	
β	Manual	B&P
0.3	1083	920
0.6	1214	761
0.9	1346	604

Table 1: Objective functions of company operation and B&P algorithm for different β values.

Table 1 shows that the model improved performance for all instances (with improvements between 15% and 45%). We notice a more significant improvement for region A,

Instance A	Manual		B&P			Instance B	Manual		B&P	
β	Δ	T	Δ	T		β	Δ	T	Δ	T
0.3	910	772	717	568		0.3	1390	952	1325	746
0.6	910	772	486	717		0.6	1390	952	594	1012
0.9	910	772	427	1065		0.9	1390	952	544	1147

Table 2: Performance of the company operation compared with the B&P algorithm in terms of total time-window violation and total travel time for different β values.

probably because access and routing are more complicated in this region. Moreover, as β increases, the improvement provided by the model becomes more significant. This is not surprising since larger values of β place more emphasis on time-window violations in the objective function. Increasing β thus yields routes in which most technician arrival times fall in a range quite close to the target response times.

Table 2 shows that in both cases (zones A and B), the results are sensitive to the value of β. This is reflected in the ratio of time-window violations to travel times in each case (namely $SH = \frac{\Delta}{T}$). Thus, in the case of zone A (B) SH goes from 1.8 to 0.5 (from 1.9 to 0.4) when β increases from 0.3 to 0.9. From the results obtained for this particular instance, we can observe how important the selection of the parameter β could become in the final performance of the model, both in terms of time-window violations and travel times.

We also carried out a sensitivity analysis with respect to two other parameters, to test not only the quality of the solutions but also the performance of the B&P algorithm. The parameters are the number of technicians (for the A instance ranging from 6 to 11 and for the B instance from 5 to 10) and the upper bound for the time window assigned to each service request with respect to the actual limit (denoted $tw1$). The idea is to run the model under tighter time-window conditions. We do this by decreasing the upper bound by 15% to 40%, thus generating new upper levels, $tw2 = 0.85 \cdot tw1$ and $tw3 = 0.6 \cdot tw1$. Table **??** gives the results of these tests for instance A with $tw1$, $tw2$, and $tw3$, while Table **??** gives the results for instance B with $tw1$, $tw2$, and $tw3$.

When we integrate regions A and B, the problem has 76 service requests (instance AB). This allows testing on a larger instance and also allows the dispatcher to consider more flexible options than those provided by a zoning fixed a priori. Eventually, the zoning could be defined by a more formal method than the traditional subdivision of the service area

by the company. Zoning may no longer be necessary (the company could completely rely on the optimization), or a clustering-type method could be used to define zones in a more systematic way. Table **??** gives the results of these experiments for 11, 13, 15, 17, 19, and 21 technicians, considering also the sensitivity with respect to the time-window upper bounds $tw1$, $tw2$, and $tw3$.

The columns of Tables **??**, **??**, and **??** represent the following: (tw) the upper bound considered for time windows, (β) the objective weighting parameter, and (#Tec) the number of technicians. For each instance generated by a combination of these parameters, we report the objective function obtained by the linear relaxation at the root node, the integer solution obtained using the default B&B implemented in CPLEX with only the columns found at the root node, and the integer solution obtained with the proposed B&P procedure (columns LR, BB, and BP respectively). We then report the gaps between (BB-LR) and (BP-LR), the number of columns generated at the root node (BB), the number of columns generated during the B&P (discounting the columns generated at the root node), the number of nodes explored by the CPLEX B&B, and the number of nodes explored by the B&P algorithm. Finally, the total running time to solve the master problem (MP) and the subproblem (SP) for both steps, namely B&B and B&P, are reported in seconds.

When the gap between the B&B solution and the linear relaxation at the root-node solution (column BB-LR) was lower than 3%, we did not run the B&P procedure. The tables report "-" for these cases.

To assign the technicians manually, the dispatchers were forced to split the region into zones A and B. Comparing the results for instances A and B in Tables **??** and **??** with those for instance AB in Table **??**, we can see that this division leads to solutions of significantly lower quality in all cases. If, for example, we consider instance ($tw1$, $\beta = 0.9$) with 6 technicians in zone A and 5 in zone B (thus, a total of 11 for AB), we see that the objective-function value obtained is 66 for case A and 1117 for case B, while that for case AB is 194. This deterioration in the objective value is due to the loss of flexibility when each zone is managed independently, since in zone B many customers need to be postponed to the next day. One conclusion to draw from these results is the considerable advantage of having an implementation of the B&P algorithm that can find an optimal solution for the larger 70-customer problem.

We allowed a maximum of 300 seconds to solve the subproblem at each iteration. We mark with a * those instances where this limit was reached before an optimal solution was found. It is clear that in these cases the bound obtained is not valid and the final solution is not optimal. As noted earlier, after solving the MIP with the columns found at the root node (B&B), if the gap between the linear relaxation at the root node and B&B was below 3%, we did not run the B&P. When used (28 of the 162 cases), the B&P algorithm reduced the gap by 3.5% on average. And the gap was reduced for more than 0.1% in 17 of the 28 cases. Most of the computational time was spent on the subproblem phase. Instance A took much longer than instance B (mainly because there are more service requests in instance A), and instance AB took the longest. Except in five cases (all from instance B) when the B&P procedure was used, fewer nodes were explored than those inspected by the B&B algorithm.

With regard to the number of columns generated, for B instances it suffices to generate around 500 columns at the root node of the B&B to solve the linear relaxation to optimality. In the larger AB instances around 2500 columns are required. When the B&P algorithm is used, the number of columns required does not seem to follow a clear pattern. Note that some values in column BP-LR (Table **??**) are negative, reflecting the difference between the B&P and the linear relaxation at the root node. These correspond to cases where the column generation at the root-node LP was terminated before reaching optimality.

Recall that we run the B&P algorithm only if the difference between LR and BB is

larger than 3%. This does not occur often, indicating that the observed gap between LR and BB is generally small. However, the solution of the linear relaxation at the root node in terms of θ is fractional in most cases.

To illustrate the behavior of the solutions, we consider two instances in Tables 3 and 4. The first case (instance 1: $tw2$, $\beta = 0.9$, $\#Tec = 6$) shows a considerable BB-LR gap (12.4%), justifying the application of the B&P procedure. In contrast, the second case (instance 2: $tw3$, $\beta = 0.3$, $\#Tec = 8$) yields a small BB-LR difference (0.3%), and the B&B algorithm obtains a close-to-optimal solution without having to explore nodes other than the root. In the tables, we show the columns considered in the linear relaxation at the root-node solution in each case, with the corresponding $\theta > 0$ value associated with each column. Although both solutions are fractional, there are two aspects that could explain

Instance 1	c_r	θ_r	$c_r \cdot \theta_r$
r1	11.0	0.828	9.1
r2	12.5	0.237	3.0
r3	11.2	0.250	2.8
r4	12.5	0.113	1.4
r5	12.0	0.363	4.4
r6	11.2	0.250	2.8
r7	12.1	0.250	3.0
r8	12.4	0.207	2.6
r9	12.7	0.500	6.4
r10	12.0	0.137	1.6
r11	11.4	0.250	2.9
r12	11.6	0.650	7.5
r13	11.0	0.172	1.9
r14	12.5	0.150	1.9
r15	11.4	0.457	5.2
r16	11.4	0.293	3.3
r17	11.6	0.250	2.9
r18	11.4	0.250	2.9
r19	14.9	0.172	2.6
other routes (~ 30)	28.4	< 0.100	2.8
		Obj. LR	70.9
		Obj. BB	79.7
	Gap [%] BB-LR		12.4

Table 3: Solution instance 1.

Instance 2	c_r	θ_r	$c_r \cdot \theta_r$
r1	73.3	0.250	18.3
r2	70.9	0.750	53.2
r3	56.0	0.500	28.0
r4	71.4	0.250	17.9
r5	70.2	0.750	52.7
r6	73.2	0.250	18.3
r7	76.3	0.250	19.1
r8	68.2	0.250	17.1
r9	63.4	0.250	15.9
r10	60.7	0.250	15.2
r11	56.2	0.250	14.1
r12	78.8	0.250	19.7
r13	80.1	0.250	20.0
r14	74.9	0.250	18.7
r15	94.7	0.750	71.0
r16	82.7	0.250	20.7
r17	79.7	0.750	59.8
r18	57.4	0.500	28.7
r19	56.7	0.250	14.2
r20	82.0	0.250	20.5
r21	84.7	0.250	21.2
r22	99.4	0.250	24.9
		Obj. LR	588.8
		Obj. BB	590.4
	Gap [%] BB-LR		0.3

Table 4: Solution instance 2.

the difference in the BB-LR gap: the nature of the fractions (θ values) and the number of columns in the solution. The LR of instance 1 (large gap) results in a large range of θ values. However, instance 2 has only three θ values (0.25, 0.5, 0.75) in all the chosen columns. This

occurred in most of the cases with a small BB-LR gap (see Tables **??**, **??**, and **??**). Moreover, the number of columns included in the linear relaxation at the root-node solution of instance 1 is much larger than the number for instance 2 (more than 50 compared with 22).

In summary, most of the solutions with a small BB-LR gap have a small number of columns in the linear relaxation at the root-node solution and a limited number of θ values (0.25, 0.5, 0.75). These examples show that when the BB-LR gap is large, the linear relaxation at the root-node solution has a range of fractional values (see instance 1) shared among a large number of columns.

6. Conclusions

In this paper, a technician dispatch problem is modeled as a vehicle routing problem with soft time windows. We solved several real examples provided by the operation of a large companyXerox Chile in different service areas of Santiago, reformulating the problem using a set covering model and using constraint programming to solve the subproblems. Constraint programming allowed the subproblems to be solved easily and effectively.

Our model allows us to optimize the dispatching of technicians in two actual instances (areas A and B as shown in Fig. 4), showing the advantages of the proposed B&P algorithm over the manual solutions implemented by the company. In addition, we were able to run our model over the entire area AB (70 service requests), obtaining in some cases significant performance improvements by not restricting the solution to the prespecified subregions A and B.

Most instances had a small gap between the linear relaxation at the root node and the B&B solution, so it was not necessary to use B&P to explore beyond the root node. When B&P was used, it reduced the gap by 3.5% on average.

References

Baldacci R., Mingozzi A., Roberti R., 2012. Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. European Journal of Operational Research 218, 1–6.

Barnhart, C., Hane, C., Vance, P., 2000. Using branch and price and cut to solve origin-destination integer multicommodity flow problems. Operations Research 48(2), 344–347.

Blakeley, F., Bozkaya, B., Cao, B. Y., Hall, W., Knolmajer, J., 2003. Optimizing periodic maintenance operations for Schindler elevator corporation. Interfaces 33, 67–79.

Bräysy O., Gendreau M., 2005a. Vehicle routing problem with time windows, part I: route construction and local search algorithms. Transportation Science 39(1), 104–118.

Bräysy O., Gendreau M., 2005b. Vehicle routing problem with time windows, part II: meta-heuristics. Transportation Science 39(1), 118–139.

Barnhart, C., Johnson, E. L., Nemhauser, G. L., Savelsbergh, M., Vance, P., 1998. Branch-and-price: Column generation for solving huge integer programs. Operations Research 3, 316–329.

Chabrier, A., 2006. Vehicle routing problem with elementary shortest path based column generation. Computers & Operations Research 33(10), 2972–2990.

Cordeau, J., Desaulniers, G., Desrosiers, J., Solomon, M., Soumis, F., 2002. The Vehicle Routing Problem. SIAM, Ch. 7: VRP with time windows, pp. 157–193.

Cordeau J., Laporte G., Pasin F., Ropke, S., 2010. Scheduling technicians and tasks in a telecommunications company. Journal of Scheduling 13, 393–409.

Dantzig, G., Wolfe, P., 1960. The decomposition algorithm for linear programming. Operations Research 8, 101–111.

Desaulniers G., Lessard F., Hadjar A., 2008. Tabu Search, Partial Elementarity, and Generalized k-Path Inequalities for the Vehicle Routing Problem with Time Windows. Transportation Science 42(3), 387–404.

Desrochers, M., Desrosiers, J., Solomon, M., 1992. A new optimization algorithm for the vehicle routing problem with time windows. Operations Research 40, 342–354.

Desrosiers, J., Solomon, M., Soumis, F., 1995. Handbooks in Operations Research and Management Science 8: Network Routing. Elsevier, Ch. 2: Time constrained routing and scheduling, pp. 35–139.

Fahle, T., Junker, U., Karisch, S., Kohl, N., Sellmann, M., Vaaben, B., 2002. Constraint programming based column generation for crew assignment. Journal of Heuristics 8(1), 59–81.

Gualandi, S., Malucelli, F., 2009. Constraint programming-based column generation. 4OR: A Quarterly Journal of Operations Research 7(2), 113–137.

Ibaraki T., Imahori S., Nonobe K., Sobue K., Uno T., Yagiura M., 2008. An iterated local search algorithm for the vehicle routing problem with convex time penalty functions. Discrete Applied Mathematics 156(11), 2050–2069.

Ilog, 2003a. ILOG CPLEX 9.0 User's Manual.

Ilog, 2003b. ILOG SOLVER 6.0 User's Manual.

Junker, U., Karisch, S., Kohl, N., Vaaben, B., 1999a. Constraint programming based column generation for crew assignment. In: International Workshop on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CP-AI-OR'99).

Junker, U., Karisch, S., Kohl, N., Vaaben, B., Fahle, T., Sellmann, M., 1999b. A framework for constraint programming based column generation. Lecture Notes in Computer Science, 261–274.

Kovacs, A., Parragh, S., Doerner, K., Hartl, R., 2012. Adaptive large neighborhood search for service technician routing and scheduling problems. Journal of Scheduling 15(5), 579–600.

Liberatore, F., Righini, G., Salani, M., 2011. A column generation algorithm for the vehicle routing problem with soft time windows. 4OR 9(1), 49-82.

Nagata Y., Bräysy O., Dullaert W., 2010. A penalty-based edge assembly memetic algorithm for the vehicle routing problem with time windows. Computers & Operations Research 37(4), 724–737.

Pillac, V., Gueret, C., Medaglia, A. L., 2013. A parallel matheuristic for the technician routing and scheduling problem. Optimization Letters 7, 1525–1535.

Pisinger D., Ropke S., 2007. A general heuristic for vehicle routing problems. Computers & Operations Research 34(8), 2403–2435.

Régin, J-C., 2011. Global constraints: A survey, in Hybrid Optimization, M. Milano and P. Van Hentenryck (eds), Springer Optimization and its Application, Vol. 45, pp. 63–134.

Rousseau, L., Gendreau, M., Pesant, G., Focacci, F., 2004. Solving VRPTWs with constraint programming based column generation. Annals of Operations Research 130(1), 199–216.

Rousseau, L., Pesant, G., Gendreau, M., 2001. Building negative reduced cost paths using constraint programming. Lecture Notes in Computer Science, 778–778.

Rousseau, L., Gendreau, M., Feillet, D., 2007. Interior point stabilization for column generation. Operations Research Letters 35(5), 660–668.

Ryan, D., Foster, B., 1981. Computer Scheduling of Public Transport: Urban Passenger Vehicle and Crew Scheduling. North-Holland, Amsterdam, Ch. An integer programming approach to scheduling, pp. 269–280.

Savelsbergh, M., Sol, M., 1991. Drive: Dynamic routing of independent vehicles. Operations Research 46(4), 474–490.

Shaw., P., 1998. Using constraint programming and local search methods to solve vehicle routing problems. Principles and Practice of Constraint Programming - Cp98 1520, 417–431, 1998

Souyris S., Cortés C. E., Ordoñez F., Weintraub A., 2013. A robust optimization approach to dispatching technicians under stochastic service times. Optimization Letters 7(7), 1549–1568.

Taillard, E., Badeau, P., Gendreau, M., Guertin, F., Potvin, J.Y., 1997. A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows. Transportation Science, 31(2),170–186.

Tang, H., Miller-Hooks, E., Tomastik, R., 2007. Scheduling technicians for planned maintenance of geographically distributed equipment. Transportation Research Part E-Logistics and Transportation Review 43, 591–609.

Vance, P., Atamturk, A., Barnhart, C., Gelman, E., Johnson, E., Krishna, A., Mahidhara, D., Nemhauser, G., Rebello, R., 1997. A heuristic branch-and-price approach for the airline crew pairing problem. Tech. rep., Georgia Tech, Supply Chain and Logistics Institute.

Vanderbeck, F., 2000. On Dantzig-Wolfe decomposition in integer programming and ways to perform branching in a branch-and-price algorithm. Operations Research 48(1), 111–128.

Vanderbeck, F., 2005. Column Generation. Kluwer, Ch. Implementing mixed integer column generation.

Vanderbeck, F., 2006. Branching in branch-and-price: A generic scheme. Working Paper, Applied Mathematics U-05.14, University Bordeaux.

Vanderbeck, F., Savelsbergh, M., 2006. A generic view of Dantzig-Wolfe decomposition in mixed integer programming. Operations Research Letters 24(3), 296–306.

Vidal T., Crainic T., Gendreau M., Prins C., 2013. A hybrid genetic algorithm with adaptive diversity management for a large class of vehicle routing problems with time-windows. Computers & Operations Research 40, 475–489.

Weigel D., Cao B., 1999. Applying GIS and OR techniques to solve Sears technician- dispatching and home-delivery problems. Interfaces 29(1), 112–130.

Weintraub, A., Abud, J., Fernandez, C., Laporte, G., Ramirez, E., 1999. An emergency vehicle dispatching system for an electric utility in Chile. Journal of the Operational Research Society 50(7), 690–696.

Xu J., Chiu, S.Y., 2001. Effective heuristic procedures for a field technician scheduling problem. Journal of Heuristics 7, 495–509.

Yunes, T., Moura, A., de Souza, C., 2000. A hybrid approach for solving large scale crew scheduling problems. Lecture Notes in Computer Science, 293–307.

Yunes, T., Moura, A., De Souza, C., 2005. Hybrid column generation approaches for urban transit crew management problems. Transportation Science 39(2), 273–288.